CPS: Beyond Usability: Applying Value Sensitive Design Based Methods to Investigate Domain Characteristics for Security for Implantable Cardiac Devices

Tamara Denning
University of Utah
tdenning@cs.utah.edu

Batya Friedman
University of Washington
batya@uw.edu

Brian Gill
Seattle Pacific University
bgill@spu.edu

Daniel B. Kramer
Beth Israel Deaconess Medical Center
dkramer@bidmc.harvard.edu

Matthew R. Reynolds
Harvard Clinical Research Institute
matthew.reynolds@hcri.harvard.edu

Tadayoshi Kohno
University of Washington
yoshi@cs.washington.edu

ABSTRACT

Wireless implantable medical devices (IMDs) are cyber-physical systems that deliver life-saving treatments to cardiac patients with dangerous heart conditions. Current access control models for these systems are insufficient; more security is necessary. In response to this problem, the technical security community has investigated new directions for improving security on these resource-constrained devices. Defenses, however, must not only be technically secure; in order to be deployable, defenses must be designed to work within the needs and constraints of their relevant application spaces. Designing for an application space—particularly a specialized one—requires a deep understanding of the stakeholders, their values, and the contexts of technology usage. Grounding our work in value sensitive design (VSD), we collaborated as an interdisciplinary team to conduct three workshops with medical providers for the purpose of gathering their values and perspectives. The structure of our workshop builds on known workshop structures within the human-computer interaction (HCI) community, and the number of participants in our workshops (N=24) is compatible with current practices for inductive, exploratory studies. We present results on: what the participants find important with respect to providing care and performing their jobs; their reactions to potential security system concepts; and their views on what security system properties should be sought or avoided due to side effects within the context of their work practice. We synthesize these results, use the results to articulate design considerations for future technical security systems, and suggest directions for further research. Our research not only provides a contribution to security research for an important class of cyber-physical systems (IMDs); it also provides an example of leveraging techniques from other communities to better explore the landscape of security designs for technologies.

ACM 978-1-4503-3005-3/14/12···$15.00
http://dx.doi.org/10.1145/2664243.2664289

Keywords

Cyber-physical systems, human factors, implantable cardioverter-defibrillators, implantable cardiac devices, implantable medical devices, medical, methods, pacemaker, practical security, privacy, security, stakeholders, envisioning workshops, value sensitive design.

1. INTRODUCTION

This work focuses on security for an important class of cyber-physical systems: wireless implantable medical devices. In this work, we bring together the technical computer security research community and value sensitive design—a framework first developed within the human-computer interaction community—in order to gather the kind of domain-specific information that is necessary to design and deploy effective security.

On the Need for Value-Sensitive Investigations. Designing effective security systems that will be appreciated and embraced by users and other stakeholders requires a deep understanding of stakeholders, their values, and the contexts of technology use. This is especially true for systems like cyber-physical systems, which perform critical functions and can require non-standard security and access control solutions. In order for the security community to design effective security that can be deployed in the real world, it is critical that human beings and human-centered methods (e.g., [10], [27]) be foregrounded in the research process.

Prior research has explored users' internal mental models in order to help security researchers design systems that act in accordance with users' expectations; this kind of work also helps security researchers understand the reasons behind unexpected user behaviors (e.g., [13][38]). Other work examines how to design for and evaluate the performance of usable security and access control systems (e.g., [2][4][5][19][33]). While this body of research is valuable, the studies are often focused solely on users, and often give little consideration to other important stakeholders or the larger ecosystem in which a technology and its security system are likely to be deployed.

As part of our research agenda—and drawing on value sensitive design—we have identified four additional elements that we believe are essential to consider in the computer security research and design process:

- **Stakeholders (e.g., [6], [14], [24],).** Who are the broad set of people who have a direct or indirect stake in the system? Indirect stakeholders are not users, but may nonetheless be impacted significantly by the technology.
- **Stakeholder Goals and Values (e.g., [6], [7], [14], [24], [41]).** What are the diverse stakeholder groups trying to accomplish, and what do they value?
- **Implications for Values (e.g., [6], [12], [14], [24], [25]).** How might different security design choices potentially impact stakeholder values—positively or negatively? In what ways might direct and indirect stakeholder groups be differently impacted?
- **Communicating Technical Concepts in Comprehensible Ways (e.g., [6], [7], [26]).** How can security researchers and designers communicate about technical security concepts with stakeholders in ways that stakeholders can understand and appreciate?

Wireless Implantable Medical Devices. Many modern implantable medical devices, including pacemakers and defibrillators, are both computational and wireless. These cyber-physical systems both contain private information and can physically affect patients' bodies. The combination of these properties makes IMDs a critical class of devices to secure.

Our Work. In this study we draw upon value sensitive design (see Section 2.2) to further these agenda elements using the domain of implantable cardiac devices—a domain that has already received significant interest from the community (e.g., [15][37][30][18]).

We conducted security-oriented Envisioning Workshops (see Section 4.2) with a variety of stakeholders involved in the care of patients with implantable cardiac devices, including: nurses, emergency physicians, cardiologists, anesthesiologists, and device manufacturer representatives. We present results on: (1) what participants find important with respect to providing care and performing their jobs; (2) the metaphors participants use to describe implantable cardiac devices and security systems for these devices; (3) participants' evaluations of potential systems that represent different directions in technical security design; and (4) participants' views on what security system properties should be sought or avoided due to domain-relevant side effects. To be clear, the purpose of this research is not to gather participants' feedback on the security performance of these systems—after all, the participants are not security experts—but rather to *gather information about how different access control systems might impact participants' jobs and their ability to care for patients*. This information helps security researchers design solutions that avoid negative side effects and tailors solutions to the needs of the domain. While much of the gathered data might be considered common knowledge by those with experience in the area, we refine and utilize procedures with which to gather this information in a structured manner.

Contributions from this work are as follows:

- **Domain Information for Designing Effective Security.** We offer domain-specific findings for implantable cardiac devices—cyber-physical systems that need improved security. Security experts can utilize the data from this study to inform the design of security systems, with the goals of increasing system adoption, supporting correct usage of security systems, and avoiding negative system side effects.
- **Method.** We adapted the Envisioning Workshop method to the security domain: we foregrounded early-stage security systems in order to gather explicit feedback on potential security directions and to identify value tensions. We provide a case study for the computer security community that demonstrates how researchers can draw upon stakeholder expertise to understand the relevant properties of a new technology domain, including value tensions among stakeholder groups. While we have conducted this research in the context of implantable medical devices, the workshop technique could be used to explore other domains such as automobiles, augmented reality, or 3D printing.

2. RELATED WORK

2.1 Technical Work in IMD Security

Implantable cardiac devices store information such as the patient's name and records of irregular heart rhythms that occurred since the last checkup. Patients visit cardiology clinics periodically so that medical staff can download information about these episodes and adjust settings on the patient's device. Current-generation cardiac devices have the ability to communicate wirelessly with external equipment from distances up to 5 meters away. There are numerous reasons for making IMDs wireless. For example, wireless IMDs can be configured ("programmed") by doctors in the operating room from farther away, which avoids the need to bring programming equipment into the sterile operating area. Wireless technology also allows the IMD to send alerts to a home monitoring station—which can then send a report to the patient's physician for analysis—without causing interruption to the patient's activities. Unfortunately, incorporating a new wireless interface for these devices also increases the communication surface on which they can be attacked, and current security models do not provide access control on a per-device basis. Here we provide an overview of prior work on security for IMDs.

In 2008, it was demonstrated that an implantable cardiac device with centimeters-range wireless communications capabilities can be wirelessly compromised by a (nearby) unauthorized party [18]; in 2011, it was demonstrated that an implantable cardiac device with meters-range wireless capabilities can also be wirelessly compromised [15]. Other work has demonstrated vulnerabilities in wireless insulin pump systems [23][28].

Challenges for securing wireless IMDs and possible directions for improving future security have been outlined [17]. One key challenge is to balance security (*blocking inappropriate access*) while also providing some guarantee that safety can be ensured in an emergency (*facilitating appropriate situational access*). To illustrate this point, consider a security system in which the IMD only grants wireless access to individuals who know a password, such as the patient's regular cardiologist. While such a system does improve security and can prevent unauthorized access by random individuals, this system also directly and negatively impacts safety: emergency personnel will not be able to read or change settings on the device without first contacting the patient's cardiologist, who might be unreachable.

There have been numerous early-stage proposals to help improve security while prioritizing medical access. One proposed direction requires the patient to wear a wristband that protects the security of the IMD when worn, but that can be removed for emergency access [8][15][1]. Another direction requires body modifications, such as RFID implants or tattoos with visible or UV-visible ink [32]. Yet another direction requires the doctor to place something on or near the patient in order to activate longer-range wireless

capabilities (e.g., [3][30][37]), taking advantage of cryptographic distance-bounding, intra-body signaling, or physiologically-derived keys. Drawing from past work, an IMD could also potentially use automated techniques to detect emergency situations and decrease security requirements (e.g., not require a password) if the patient is in a state of medical emergency [16]. Another, more traditional approach might be to issue temporary or permanent access passwords via a centralized entity, such as a manufacturer-maintained database.

Although less related to our work on implantable cardiac devices, there has been significant work focused on security and privacy for personal medical sensors and networked medical devices (e.g., [29]); see [1] for a survey. Many of the efforts in this space also have potential applicability to implantable medical devices. For example, the Amulet system [34]—which requires the user to wear an external device—has many overlapping elements with other defenses for IMDs [8][15][1]. Additionally, any effort to improve key establishment for body-area networks (e.g., [3][37]), can be used to help improve the security of key establishment systems for IMDs.

Recently, a systematization of knowledge paper was published on the topic of computer security for implantable medical devices and body area networks (BANs) [31]. The paper provides a thorough examination of both attack and defense work dealing with these classes of technologies in the computer security community. The authors categorize defense directions as falling into four different trends: Biometric and Physiological Values (e.g., ECG or IPI), Out-of-Band (e.g., tattoos), Distance Bounding (e.g., intrabody signaling or cryptographic), and External Devices (e.g., fail-open wristband). Four additional categories—Wireless Attacks, Software/Malware, Anomaly Detection, and Emerging Threats—are used to classify other research trends in the area.

2.2 Value Sensitive Design

Computer security and access control systems are frequently discussed in the context of values such as security, privacy, and convenience. These systems, however, also affect and are affected by other important human values such as trust, physical welfare, autonomy, or human dignity. In this research we drew on established theory and methods from value sensitive design (e.g., [12][14][24][26]) to frame our study design and our data analyses. First developed in human-computer interaction,, value sensitive design has since been used in civil engineering, information management, human-robotic interaction, and ubiquitous computing. For example, one security-focused series of value sensitive design studies analyzed informed consent for cookies in web browser security [25], leading to recommendations for browser redesign to better support informed consent and a proof-of-concept redesign in the form of a plug-in "cookie watcher" for the Mozilla browser [12]. Another security-focused study investigated users' mental models for web browser security, suggesting that elements in the user interface (e.g., the open or closed padlock) were inadvertently leading some users to construct incorrect mental models for a secure connection [13]. More recently, value sensitive design methods have been applied in research about safety and security for mobile phone parenting technologies for teens [6], home technologies [9], and implantable cardiac devices from patients' perspectives ([7], see Section 3.1).

In this study, we drew explicitly on two value sensitive design methods: direct and indirect stakeholder analyses, and value dams and flows.

Direct and Indirect Stakeholder Analyses. In examining the ecosystem surrounding security for implantable cardiac devices, an important question is what professional roles should be represented among the study participants. Value sensitive design stresses the consideration of both: the direct stakeholders who will interact with the technology (e.g., cardiologists); and indirect stakeholders who—while they do not directly interact with the technology—can affect and be affected by the technology (e.g., venture capitalists who invest in implantable cardiac devices). Section 3.2 provides a brief explanation as to why taking this broader view is valuable.

Value Dams and Flows. Given a wide range of possible technical security, it is not always obvious how to choose which system to pursue. Value dams and flows is a technique for identifying reasonable, value-sensitive design options from among a set of possible designs or technical features (e.g., [6][7][24]). First, options that a threshold percentage of stakeholders strongly object to are removed from the list of viable solutions (value dams); then, from the remaining options, those that many stakeholders favor are selected as good candidates for solutions (value flows). In this research we use the value dams and flows method to help identify viable security designs for implantable cardiac devices.

3. STAKEHOLDERS, SECURITY AND THE MEDICAL ECOSYSTEM FOR IMDS

3.1 Patients: Prior Work

Previous work explored patient views and values regarding their implantable cardiac devices in order to inform the design of security systems for wireless IMDs [7]. The findings in that study—as well as the absence of the medical provider viewpoint—are part of the motivation for this study.

In the patient study, semi-structured interviews were conducted with 13 individuals with implanted cardiac devices. Key questions concerned the evaluation of eight mockups of early-stage IMD security systems. For each system, patients provided judgments as to whether they liked or disliked the system and whether or not they would choose to use it. Qualitatively, patients articulated a number of values that affected their attitudes toward the systems, including: security; safety; privacy; aesthetics; psychological welfare; convenience; cultural and historical associations; self-image and public persona; and autonomy and notification. Importantly, different patients identified different sets of cares and concerns that, in turn, led to different levels of satisfaction with the proposed security designs.

3.2 Medical Providers

A patient's medical care is an ongoing process that is affected by regulation, device manufacturers, federal testing, insurance companies, hospital equipment purchases, primary care staff, specialist nurses and doctors, emergency care staff, operating room staff, and others. While the patient study described above provides some insight into the values and priorities of patients who live with implantable cardiac devices embedded in their bodies, the question of how medical providers might interact with these technical computer security directions was not addressed. Yet, to be effective, security must work for and with all of the key stakeholders. In the case of implantable cardiac devices, this includes not only the patients, but also the medical providers who—in one way or another—ensure that the devices function properly and improve patient health.

The current study with medical providers builds on the patient study by investigating similar security system designs concepts. However, as described in the methods below, the participant pool and study format differ (as appropriate to these stakeholder groups).

4. METHODS

This study is part of an on-going, interdisciplinary collaboration among researchers in computer security, human-computer interaction, and cardiology. The researchers have no special relationships with any particular medical device provider or other conflicting agency.

4.1 Participants

In this study, we sought to investigate in detail the values, priorities, constraints, and themes that emerge in a complex domain. We followed an established approach to work in depth with a smaller number of participants (e.g., [22][36]) rather than seeking to answer specific questions using a larger number of participants. We conducted three workshops with medical providers in the United States: one in a city on the west coast (Group I) and two in a city on the east coast (Groups II and III). Participants were recruited through a snowball method. The research team first sent emails to previous contacts in the medical community requesting suggestions for potential participants and relevant mailing lists; the researchers then followed up on those suggestions with email letters of invitation to participate in the research. Recruitment efforts were initially extremely slow; this was partially because we needed to obtain permission from appropriate authorities (i.e., "gatekeepers") and partially because we needed domain insiders to explain the importance of—and cultivate enthusiasm for—study participation (i.e., "advocates").

We applied for and obtained approval from our institution's human subjects review board. In order to synchronize study protocols across the multiple institutions involved in this study, it was necessary to submit multiple modifications. Participants were compensated $200 for their time; while this amount may seem unusually high, it was deemed appropriate in the context of the particular participant pool (e.g., cardiologists).

A total of 24 medical providers (age: average=39, min=28, max=64) participated in the study. Table 1 breaks down participant gender by workshop. Participants had a broad spectrum of roles in the medical ecosystem: cardiologists and electrophysiologists (n=2), nurses and nurse practitioners in cardiology and electrophysiology (n=5), anesthesiologist (n=1), emergency physician (n=1), other physicians (n=2), physician assistant (n=1), medical residents (n=4), medical device manufacturer representative (n=1), biomedical informatics researcher (n=1), and venture capitalist (n=1).

Group	Male	Female	**Total**
I	4	6	**10**
II	4	3	**7**
III	5	2	**7**
TOTAL	**13**	**11**	**24**

Table 1. Number of participants by group and gender.

4.2 Workshop Format

To elicit participant values, priorities, and constraints for the security of implantable cardiac devices, we sought a method that would provide opportunities for open-ended ideation about device security as well as focused reactions to potential early-stage security concepts. We drew inspiration from and adapted Kensing and Madsen's techniques for "generating visions" [20]—which integrates metaphorical design with a Future Workshop (particularly the critique phase)—and Yoo et al.'s Envisioning Workshop, which emphasizes surfacing value tensions among diverse stakeholders [40]. In both instances, we tailored the workshop structure to focus on security aspects of specific system designs. In addition, we sought both to collect individual reflections and to benefit from group discussion; thus, data collection included individual written materials as well as verbal group interactions. We describe the workshop protocol below.

Each session lasted a total of two hours. Audio recordings were made of each session and then later transcribed for analysis.

Implantable Cardiac Device Overview and Initial Perspectives. To ensure that all participants had some shared vocabulary for implantable cardiac devices, a research team member provided a brief overview of implantable cardiac devices and clarified how terms would be used during the workshop. This overview did not include information on security for implantable cardiac devices. Following this overview and to tap into participants' perspectives prior to any influence from the workshop activities, participants were asked to complete a brief paper and pencil worksheet that elicited their initial views on security and access control for implantable cardiac devices. The worksheet contained the following questions: (1) *What properties about implantable cardiac devices or the ecosystem surrounding them do you value most?*; (2) *What things about implantable cardiac devices or the ecosystem surrounding them should not change?;* (3) *What things about implantable cardiac devices or the ecosystem surrounding them most need improvement?*; (4) *What is the most common problem related to implantable cardiac devices that you encounter in your line of work (e.g., lack of access to patient information, inability to access cardiac device, device malfunction)?*; and (5) *What is the problem with the most negative health impact (related to implantable cardiac devices) that you encounter in your line of work?*

Metaphor Generation. To help understand the broad backdrop of participants' perspectives as well as potential mental models, participants were invited as a group to share verbally: (1) metaphors for implantable cardiac devices; and (2) metaphors for security and access control for those devices. A research team member facilitated the contributions and recorded each metaphor in a few concise words on a whiteboard.

Critiques and Concerns. To understand how security and access control systems for implantable cardiac devices could go awry as well as to understand medical providers' hesitations and concerns about this type of technology, participants were invited as a group to share verbally their concerns and fears about security for implantable cardiac devices. Volunteer participants grouped the concerns into clusters based on similarity.

Evaluation of Security and Access Control System Concepts. To understand participants' views on what properties to advocate for and which to avoid in the development of security and access control solutions for implantable cardiac devices, a researcher introduced participants to six potential security and access control systems, one at a time. The researcher indicated that these were early, representative systems designed to elicit feedback. The system concepts are described in Section 4.3 below. Directly after each system presentation, participants completed a paper and pencil worksheet in which they recorded their responses to the following questions: *From your perspective as a professional who*

A. Medical Alert Bracelet with Password

B. Centralized Database

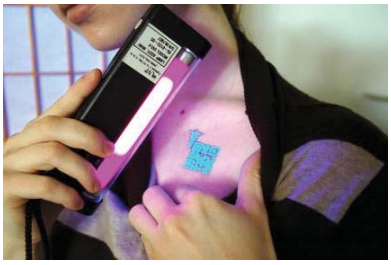

C. UV-Visible Tattoo of a Password

D. Fail-Open/Safety Wristband

E. Proximity-Based Equipment

F. Criticality-Aware Fail-Open IMD

Figure 1. Photos of the security system concepts presented to medical providers during the workshop.

deals with implantable cardiac devices, what do you like about this concept? What do you dislike about this concept? Why?

Once participants had been introduced to all six system concepts, participants completed a worksheet with the following questions: *(1) Would you say that you like any of the concepts, and if so, which ones?; (2) Would you say that you dislike any of the concepts, and if so, which ones?; (3) If you were to choose one or more of these concepts to recommend for use in the future, which concept or concepts would you choose? Why?; and (4) If you were to choose one or more of these concepts to recommend against use in the future, which concept or concepts would you choose? Why?* Participants had materials for each system concept on hand if they needed to refresh their memories.

Security System Property	System Concepts					
	A	B	C	D	E	F
Requires patient to wear something [8][15][1]	■			■		
Requires modification to the patient's body [32]			■			
Requires patient maintenance [8][15][1]				■		
Visible on the patient [8][15][1]	■		▒	■		
Access depends on centralized infrastructure		■				
Requires specialized provider equipment [3][30][37]			■		■	
Access requires proximity to the patient [3][30][37]	▒		■	▒	■	
Access has a manual override [8][15][1]				▒	■	
Security decisions are automated [16]						■

Table 2. Relevant properties of the security system concepts presented to medical providers (by system concept). Dark cells indicate a property represented by a system. Lighter cells indicate a property represented in some situations or by some interpretations.

Open-ended Discussion. Finally, to ensure that participants had ample opportunity to surface any major issues that might have been missed, participants engaged in an open-ended discussion around security and access control for implantable cardiac devices in which they could respond to and debate each other's ideas. To initiate the discussion, a research team member used the following prompt: *What are the challenges in this space?*

Due to space constraints, we do not present the results of the critiques section or the open-ended discussion. While this data contained some interesting information, by and large the content and issues overlapped with that from the other workshop sections.

4.3 Security System Concepts

As noted above, participants were asked to evaluate six security and access control system concepts (see Figure 1). A researcher presented the systems to participants using verbal explanations and slides. These systems are not complete or perfect from a security (access control: false positive), safety (access control: false negative), or usability standpoint. We presented these systems for feedback because: (a) they are representative of some of the security solutions that have been previously proposed by the security research community; (b) they represent a variety of relevant system properties; and (c) the discussion of a specific system can serve to ground an otherwise abstract discussion. Table 2 provides a summary of some of the systems' relevant properties. The six security concepts are as follows:

A. **Medical Alert Bracelet with Password.** A medical alert bracelet worn by the patient with a password engraved on the inside. The password is required to read information off of or to reprogram the implanted cardiac device via the programmer.

B. **Centralized Database.** A centralized database system that medical centers can access to obtain a temporary password that can then be used to access the implanted cardiac device.

C. **UV-Visible Tattoo of a Password.** The patient is tattooed with a scannable representation of a password (e.g., 2D barcode) using UV-visible ink that is only visible under a UV (black light) light source.

D. **Fail-Open/Safety Wristband.** The patient wears a battery-powered system that, when present, prevents unauthorized programmers from communicating with the implanted cardiac device. When the wristband is removed, the implanted device accepts communications from all programmers. Additionally, the wristband: sounds an alarm when the patient is approaching a strong magnet source (a potentially serious risk to cardiac device patients); and dials 911 if the patient is experiencing a cardiac emergency.

E. **Proximity-Based Equipment.** Equipment that medical professionals can hold in contact with the patient's body in order to communicate with the implantable cardiac device.

F. **Criticality-Aware Fail-Open IMD.** An implantable cardiac device that uses available information—such as the patient's cardiac rhythm, movements, and location—to determine whether or not the patient is likely experiencing a medical emergency. If the system detects a likely emergency, it changes the security settings to allow access from any programmer. The system does not notify anyone about the medical emergency and does not change cardiac therapies.

4.4 Coding and Reliability

Participants' written initial perspectives were coded systematically using the following process: one researcher developed a coding scheme using all of the data; once completed, that researcher used the finalized coding scheme to systematically recode the entire data set. A second coder—not affiliated with the research team or the study—was trained in the coding scheme using data from 4 participants, and then independently performed reliability coding of the data for the remaining 20 participants. This process resulted in an overall kappa of 0.745; Fleiss rates any value of kappa over 0.75 as excellent agreement and between 0.40 and 0.75 as intermediate to good agreement [11], while Landis and Koch rate a kappa of 0.81 to 1.00 as "almost perfect" and between 0.61 and 0.80 as "substantial" agreement [21].

The metaphor data set was smaller and, thus, more appropriately coded by consensus. We used the following process: (1) first, two researchers independently read through all of the data to generate an initial set of coding categories and assign responses to categories; (2) next, using consensus, researchers iteratively synthesized categories and arrived at agreement; and (3) both researchers made another independent pass through all of the data and any lingering disagreements were resolved.

Justifications in the security system concept evaluation data were identified from inspection of the qualitative data and are presented via participant quotes.

5. RESULTS

Given the relatively small number of participants in each workshop, there was no way to draw meaningful comparisons among the workshops' participant demographics (e.g., location, gender, age, professional role). We combined the data from all three workshops into one data set.

5.1 Initial Perspectives

Participants' written responses to the Initial Perspectives questions provide a relatively unbiased (that is, largely uninfluenced by our subsequent workshop activities) view into what participants consider important about implantable cardiac devices and their usage to treat patients. Since our primary interest was to understand broadly the pre-existing issues important to medical providers, we examined providers' responses to the set of five questions as a whole (rather than by individual question).

Thirteen categories of issues emerged from the analysis of participant responses as follows (in alphabetic order): (1) *Access & Sharing;* (2) *Compatibility;* (3) *Correct Usage;* (4) *Device Battery Life;* (5) *Device Compactness / Inertness;* (6) *Device Ecosystem;* (7) *Device Functionality;* (8) *Patient / Patient Health;* (9) *Programming;* (10) *Quality of Data;* (11) *Remote Monitoring;* (12) *Security & Privacy; and* (13) *Surgery & Healing.* Over three-quarters of the participants expressed issues related to Device Functionality (79%) and Patient/Patient Health (75%); more than half mentioned Surgery & Healing (58%). The next most represented categories were mentioned by roughly a quarter of the participants (ranging from 25-29%). That said, given the sample size and exploratory nature of this study, we believe it would be prudent to consider all 13 categories of issues when designing a security system for implantable cardiac devices.

This list of issues provides a window into the values and priorities of medical providers dealing with implantable cardiac devices. Security and human-computer interaction researchers may not have sufficient domain knowledge to make direct judgments as to how a system design might interact with these aspects of medical care; these categories, and other data like them, may serve as a meaningful starting point for dialog with domain experts.

5.2 Metaphors

Metaphors often underlay people's mental models of technological systems, which can affect the ways in which they interact with those systems. The metaphors supplied by participants provide some indication as to how they conceptualize implantable cardiac devices and security systems for those devices. In addition, using metaphor generation as an opening activity was intended to help break the ice; metaphor generation is rapid and appropriate for ideas that might otherwise be considered offbeat or silly.

As a group, participants generated a total of 81 metaphors: 42 for the implantable cardiac devices and 39 for security and access control for those devices. The following 11 categories—given in alphabetic order—emerged from clustering together similar metaphors: (1) *Agency;* (2) *Bio-medical;* (3) *Business;* (4) *Emotion;* (5) *Information;* (6) *Maintenance;* (7) *Personal Identity;* (8) *Privacy;* (9) *Risk;* (10) *Security; and* (11) *Technology.* Participants understood implantable cardiac devices in a broad variety of ways, including bio-medical terms, both positive (e.g., "life savers") and negative ("site of infection"); and emotional terms, though always negatively (e.g., "anxiety producing," "source of hassle"). Participants described the device's security systems in terms of: security, both secure (e.g., "secure site on the Internet") and insecure (e.g., "bank with an unlocked vault"); risk (e.g., "life threatening"); and information,(e.g., "complete control of information," "black box on a plane").

The diversity of metaphors as well the potential for any given metaphor type to convey both positive and negative dimensions points to the need for security researchers to attend carefully to how stakeholders conceptualize—in lay terms—security for such devices, and how they use those conceptualizations to generate positive or negative perspectives on the security system.

5.3 Security System Concepts

To understand participants' feedback on the six specific security system concepts presented, we adapted a values dams and flows approach. We looked at which systems participants found strongly acceptable (flows)—that is, the systems which many participants liked and very few participants disliked—and which systems participants found less acceptable (dams)—that is, the systems which few participants liked and many participants disliked. This analysis helps inform the interpretation of the following section, in which we present some of the reasons that providers expressed for liking or disliking specific systems.

Evaluations. Table 3 provides an overview of the results from participants' evaluations of the security system concepts presented. The fail-open/safety wristband (shown in Table 3 with green cells) was the best received in all categories: the largest percentage of providers liked it (58%) and would recommend its use (46%), and the smallest percentage of providers disliked it (17%) and would recommend against its use (13%). The UV-visible tattoo of a password (shown in Table 3 with the row of red cells) was the least satisfactory in all categories: only 17% of providers liked it; 54% disliked it; 50% recommended against its use; and only 13% recommended its use. Two other systems reach relatively high thresholds on dislike and recommend against: the medical alert bracelet with a password (46% dislike and 33% recommend against) and the criticality-aware fail-open IMD (42% dislike and 38% recommend against); these data might suggest avoiding the use of those three system concepts.

In general, when we examine these evaluation results, we require high satisfaction thresholds (i.e., high "like" and "recommend" percentages, low "dislike" and "recommend against" percentages) in order to describe a system as well-liked. In contrast, less stringent thresholds are necessary to describe a system as problematic, in order to respect the perspectives of stakeholders who may be in the minority. This is in line with the value sensitive design dams and flows analysis guidelines [24].

Justifications. As noted above, the evaluation results on the security system concepts immediately raise the question of why providers like or dislike a system or would recommend for or against its use. Recall that Table 2 provides a breakdown of some of the properties embodied by the various system concepts, such as requiring physical proximity to the patient or having a manual override. Providers are potentially reacting to these properties in their evaluations. Below we report what system properties providers said they liked and disliked about each system; for the systems that were particularly high-ranked or low-ranked, we break out the relevant properties as lists. Note that participants provided the justifications for each system concept directly after it was presented; participant-supplied justifications are potentially influenced by this fact. The quantitative evaluations were completed after all system concepts had been presented.

Providers N=24	Participant Percentage			
	Like	Dislike	Rec.	Rec. Against
A. Medical Alert Bracelet w/ Password	29	46	21	33
B. Centralized Database	38	21	25	25
C. UV-Visible Tattoo of a Password	17	54	13	50
D. Fail-Open/Safety Wristband	58	17	46	13
E. Proximity-Based Equipment	38	25	38	21
F. Criticality-Aware Fail-Open IMD	38	42	33	38

Table 3. Percentage of participants by security system concept who liked, disliked, recommended, or recommended against each system concept. Green indicates high satisfaction with a system concept; red indicates low satisfaction.

Medical Alert Bracelet with Password (System A). The medical alert bracelet was one of the system concepts most disliked (46%) and recommended against (33%). Providers most frequently expressed disliking System A for the following reasons:

- ↓ Access is not guaranteed—the bracelet may be forgotten, lost, stolen, damaged, or the patient may choose not to wear it. (E.g., "In an accident, the bracelet could be damaged/lost and emergency personnel would not be able to access device")
- ↓ The security is insufficient. (E.g., "Way too easy to maliciously steal password")
- ↓ It visibly indicates to the patient and others that the patient has a condition. (E.g., "Identifies pt. as 'having a problem'")

The relatively poor reception of System A suggests that either these properties are particularly disagreeable to participants, or that the advantages are not sufficient incentive to tolerate the disadvantages. When participants expressed liking that the medical bracelet solution they noted that the system did not depend on other equipment or systems, provided "reassurance" to the patient, was cheap, and provided some security.

Centralized Database (System B). The centralized database was neither one of the highest-rated nor one of the lowest-rated systems. Participants expressed concerns about: the availability of the database across regions, across providers and manufacturers, in case of disaster, or in case of other technical difficulties; how to identify patients to look them up in the database; how to secure the database and identify who is authorized to access it; who will administer the database and how they will fund and maintain it; and the fact that a database would require time away from the bedside to access. In contrast, participants appreciated that this system neither required nor depended upon the patient to wear anything, was theoretically universal, and provided more security than System A (the Medical Alert Bracelet with Password).

UV-Visible Tattoo of a Password (System C). The UV-Visible Tattoo was the lowest-ranked system for all evaluation questions. Participants expressed disliking this system for the following reasons:

- ↓ Required equipment may not be working, accessible, or timely to acquire. (E.g., "requires UV light (i.e. working bulb, power source)")
- ↓ Patients may have cultural, social, or personal objections over a tattoo. (E.g., "religious restrictions against tattoos")
- ↓ Access is not guaranteed—the tattoo could be faded, damaged, or distorted. (E.g., "blood or trauma may obscure tattoo")

↓ Password revocation or changes could be complicated. (E.g., "how to change when device is changed out")

Again, this suggests that the disadvantages outweigh the properties about the system which participants liked: its invisibility in the patient's daily life, both for human and security reasons; and the fact that it is (theoretically) always with the patient, but requires no patient effort.

Fail-Open/Safety Wristband (System D). The fail-open/safety wristband was the highest-rated system across all categories. Participants reported liking the system for the following reasons:

↑ The fail-safe mode guarantees access. (E.g., "GREAT failsafe mode (remove bracelet)")

↑ The system provides some safety features. (E.g., "safety features BIG plus")

↑ The system provides some security features. (E.g., "Provides mechanism against snoops…equivalent to locking your door when you leave the house")

↑ The mechanism gives access control power to the patient. (E.g., "pt. feels empowered. pt. is an active participant in their own care")

↑ Provides a visual cue to EMTs. (E.g., "identifies patient as having an ICD")

Following previous lines of reasoning, these advantages must outweigh the dislikes expressed by participants: that there is no security if the wristband is not worn (and it is easily removed); that the wristband requires battery replacements or recharging and requires the patient to wear it, for which there is no incentive; that there may be many false-positive calls to 911; that the patient is visibly identified as having a medical condition; that emergency medical staff would require training to know to remove the wristband; and that the system is potentially expensive to develop and produce.

Proximity-Based Equipment (System E). The proximity-based equipment was neither one of the highest-rated nor one of the lowest-rated systems; this system most closely resembles the status quo of access control for implantable cardiac devices. Participants expressed liking a variety of system properties: that it provides some security from wireless tampering; that it does not require the patient to wear or do anything (and therefore does not provide a visual indication of the patient's condition); that it does not depend upon other equipment or systems; that it is similar to the current model; that it is easy, and allows bedside access; that it would be a (theoretically) universal access system; and that it gives the patient some control over who may access their device. Conversely, participants reported disliking: that patients are still susceptible to in-person security breaches; that such a system would require new equipment, which is expensive; that such a system would potentially be manufacturer-specific; and that such a system would require all medical centers to have the equipment on-hand and readily accessible for emergencies.

Criticality-Aware Fail-Open IMD (System F). The criticality-aware IMD was one of the systems most disliked (42%) and recommended against (38%). Providers most frequently expressed disliking System F for the following reasons:

↓ The IMD may not correctly identify a medical emergency (false negative—closed access). (E.g., "this assumes the device can properly recognize emergencies → current devices can't even recognize some arrhythmias correctly")

↓ The IMD may incorrectly identify a medical emergency (false positive—open access). (E.g., "possibility of misidentifying a 'medical emergency'")

↓ There may be non-emergency situations where the IMD needs to be accessed. (E.g., "what happens if the patient moves and has a new cardiologist?")

↓ This system could change IMD size or shape, consume battery life, or cost more. (E.g., "will certainly add expense to cost of device such that CMS may veto payment/reimbursement")

As previously reasoned, these disliked properties apparently overpower the properties that participants liked: that it (theoretically) allows access in an emergency; that it provides some security; that it depends upon no extra equipment; and that it does not require the patient to do or wear anything.

6. PROVIDER VS. PATIENT RESULTS

Recall that this medical provider study was conducted in part to complement the prior study on patients' views and values regarding early-stage security solutions for implantable cardiac devices [7]. The set of systems presented to medical providers was the same as those presented to patients with the following exceptions: three somewhat redundant system concepts were removed (i.e., visible tattoo of a password, fail-open wristband, and fail-open/patient-specified-functionality wristband) and one new security concept was included: a centralized database. Additionally, the presentation format was tailored to the participants. For the medical providers, the system concepts were presented in a group setting via a verbal description and supporting slides. For the patients—many of whom were older or ill—the system concepts were presented to each individual with a verbal description and a physical mockup as a prop; patients provided verbal feedback as part of their semi-structured (individual) interviews. The questions asked, while similar in substance, were also slightly different as appropriate to the individual's role (as a medical provider or a patient).

We now consider the quantitative results from both studies together, as a way to explore security concepts that might be successful for both sets of stakeholder groups. For the proximity-based equipment approach (System E): in the patient study this security concept was least disliked (0% dislike), and hence might be the most logical system to choose; however, 25% of the medical providers disliked the proximity-based equipment approach and 17% would recommend against its use, making it a less desirable choice overall. In a similar vein, the criticality-aware fail-open IMD (System F) was more liked (27%) than disliked (18%) by patients; however it was more disliked (42%) than liked (38%) by providers. In the case of the criticality-aware fail-open system, we suspect that this difference is primarily due to providers' higher concern regarding the lack of a manual override if the system fails to recognize a medical emergency. The provider results suggest that this approach may not be a suitable solution for securing IMDs.

The patient study recommended a set of three solution choices that, if offered together, might satisfy the desires of most patients: a proximity-based system (System E), a fail-open/safety wristband (System D), and a UV-visible tattoo of a password (System C). Given the strong opposition to UV-visible tattoos among providers (50% would recommend against), however, our new results caution against their use in practice).

In terms of similarity of perspective, the fail-open/safety wristband approach (System D) was the security concept that was least disliked (17% dislike, 13% recommend against) and the most liked by medical providers (58% like, 46% recommend). This concept was also the most liked by the patients (45% like, 27% would choose to use it if it were available).

7. SECURITY DESIGN CONSIDERATIONS

Drawing on a synthesis of the results from the study data, we now provide concrete design considerations for security researchers working in the implantable cardiac device domain.

7.1 Access, Access, Access

Access—and the related issue of compatibility—show up in both the initial perspectives and the security system concept data sets, and are particularly emphasized in the latter. Participants repeatedly indicated the importance of (different kinds of) access: (a) Providers must always be able to access the implanted device, and security systems should either fail to an open state or offer some kind of override; (b) "Unplanned" access does not only occur in emergencies, since patients may travel or change cardiologists, and records are not always transferred smoothly; (c) Access should not rely upon a centralized system, which could be unavailable (due to technical, geographic, or other reasons) and which merely defers the security problem; (d) Access cannot rely upon a conscious or compliant patient; (e) Access should avoid relying on additional equipment, which can delay or block patient care or remove providers from the bedside; (f) Access should be timely, and should therefore require few steps. Perhaps, above all, in the words of one of our participants: "Please, please, please keep it SIMPLE."

7.2 Mechanics and Logistics

Various aspects of IMD mechanics and logistics are raised in both the initial perspectives and the system concept data sets. Any security system should avoid disturbing the status quo in terms of: (a) cost, which can also affect insurance approval; (b) required training, particularly for non-cardiology staff; (c) implant battery usage; (d) implant size; or (e) any other aspect that might impact the surgical or healing processes.

7.3 Safety Features and Incentives

Participants showed interest in the possibility of incorporating safety features into a security system for IMDs. The exact nature of such features and how they might be tuned should be further investigated; for example, many participants expressed concern that a 911 feature would result in many false-positive emergency notifications. However, if well designed, safety features might function as incentives for patients to comply with using a security system. This is particularly relevant in the case of a system like the wristband (System D), which incorporates an external battery which can be used for safety features, and with which the patient only receives security if they choose to wear the band.

7.4 Empowering Without Burdening

Ideally, patients should be given some implicit or explicit role in the access control process, whether via overt action or by allowing someone extended skin contact. Generally speaking, such a role might give patients a feeling of empowerment, but more practically speaking, patients could provide human reasoning as to whether or not their device should be accessed in a given situation. Conversely, patients should not be unduly mentally or physically burdened by a security system. As one example of this, anything that visually indicates the patient's condition should be opt-in; visual indicators such as medical alert bracelets are useful to emergency staff, but patients should be able to weigh the advantages and disadvantages and choose whether or not to participate. Moreover, this consideration raises a host of ethical, legal, and philosophical questions: Should a security design hinge upon patients being able to choose whether or not they wish to comply? How many patients would actually comply? Should a security design strive to equally protect all patients from potential harm, or is that attitude paternalistic? What are the repercussions, legally or in terms of reputation, if a company's IMD is attacked, and security was optional? The domain is full of interesting questions that are ripe for consideration.

8. CONCLUSION

The work reported here makes two important contributions. First, we offer domain-specific findings for security for implantable cardiac devices—cyber-physical systems that perform critical operations within patients' bodies. The purpose of these findings is to facilitate designs with increased adoption, correct usage, and few negative side effects for patients and medical providers.

Our second contribution concerns method. Specifically, we adapted established methods from value sensitive design—the Envisioning Workshop and values dams and flows—to the security domain: we gathered feedback on classes of security systems via concrete descriptions of systems that embody particular properties (security and non-security). This work demonstrates how security researchers can draw upon direct and indirect stakeholders to understand the relevant properties of a new technology domain; the techniques we used were in the context of implantable cardiac devices, but could be used to explore other emerging technology domains.

Bridging the gap between technological innovation and the lives of stakeholders who will be impacted by that technology is not easy; however, it is critical to do so. Toward that end, the work reported here provides one specific study to suggest how such work could be done, and methods for making progress toward incorporating human values into technical security design.

9. ACKNOWLEDGMENTS

This work was supported in part by the US National Science Foundation (NSF) under awards CNS-0846065, CNS-0905118, and CNS-0905384, by the US Department of Health & Human Services (HHS) under award HHS-2010-03958-09, and by an Alfred P. Sloan Research Fellowship. Any opinions, findings, and conclusions or recommendations expressed in this publication do not necessarily reflect the views of the funding agencies. We thank the medical providers who participated in this study. We thank M. Enev, D. Halperin, K. Koscher, B. Ransford, and A. Takakuwa for contributing to the data analysis process. This work was done while Tamara Denning was at the University of Washington.

REFERENCES

[1] S. Avancha, A. Baxi, and D. Kotz. Privacy in Mobile Technology for Personal Healthcare. ACM Computing Surveys, 45(1), 2013.

[2] D. Balfanz. Usable Access Control for the World Wide Web. ACSAC 2003.

[3] S. Cherukuri, K. Venkatasubramanian, and S. Gupta. BioSec: A Biometric Based Approach for Securing Communication in Wireless Networks of Biosensors Implanted in the Human Body. ICPP Workshops 2003.

[4] S. Clark, T. Goodspeed, P. Metzger, Z. Wasserman, K. Xu, and M. Blaze. Why (Special Agent) Johnny (Still) Can't Encrypt: A Security Analysis of the APCO Project 25 Two-Way Radio System. USENIX 2011.

[5] L. Cranor and S. Garfinkel. Security and Usability. O'Reilly Media, Inc., 2005.

[6] A. Czeskis, I. Dermendjieva, H. Yapit, A. Borning, B. Friedman, B.T. Gill, T. and T. Kohno. Parenting From the Pocket: Value Tensions and Technical Directions for Secure and Private Parent-Teen Mobile Safety. SOUPS 2010.

[7] T. Denning, A. Borning, B. Friedman, B. T. Gill, T. Kohno, and W. H. Maisel. Patients, Pacemakers, and Implantable Defibrillators: Human Values and Security for Wireless Implantable Medical Devices. CHI 2010.

[8] T. Denning, K. Fu, and T. Kohno. Absence Makes the Heart Grow Fonder: New Directions for Implantable Medical Device Security. HotSec 2008.

[9] T. Denning, T. Kohno, H. M. and Levy. Computer Security and the Modern Home. Communications of the ACM, 56(1), 94-103, 2013.

[10] T. Dimkov, A. van Cleeff, W. Pieters, and P. Hartel. Two Methodologies For Physical Penetration Testing Using Social Engineering. ACSAC 2010.

[11] J. L. Fleiss, B. Levin, and M. C. Paik. Statistical Methods for Rates and Proportions (3rd ed.). New York: John Wiley & Sons. 2003.

[12] B. Friedman, D. Howe, and E. Felten. Informed Consent in the Mozilla Browser: Implementing Value Sensitive Design. HICSS 2002.

[13] B. Friedman, D. Hurley, D. C. Howe, E. Felten, and H. Nissenbaum. Users' Conceptions of Web Security: A Comparative Study. CHI 2002 (Extended Abstracts).

[14] B. Friedman, P. H. Kahn Jr., and A. Borning. Value Sensitive Design and Information Systems: Three Case Studies. In P. Zhang and D. Galletta, editors, Human-Computer Interaction and Management Information Systems: Foundations. 2006.

[15] S. Gollakota, H. Hassanieh, B. Ransford, D. Katabi and K. Fu. They Can Hear Your Heartbeats: Non-Invasive Security for Implanted Medical Devices. ACM SIGCOMM 2011.

[16] S. K. S. Gupta, T. Mukherjee, and K. Venkatasubramanian. Criticality Aware Access Control Model for Pervasive Applications. IEEE PERCOM 2006.

[17] D. Halperin, T. S. Heydt-Benjamin, K. Fu, T. Kohno, and W. H. Maisel. Security and Privacy for Implantable Medical Devices. IEEE Pervasive Computing, 7, 2008.

[18] D. Halperin, T. S. Heydt-Benjamin, B. Ransford, S. S. Clark, B. Defend, W. Morgan, K. Fu, T. Kohno, and W. H. Maisel. Pacemakers and Implantable Cardiac Defibrillators: Software Radio Attacks and Zero-Power Defenses. IEEE S&P Symposium 2008.

[19] A. Jøsang, B. AlFayyadh, T. Grandison, M. AlZomai, and J. McNamara. Security Usability Principles for Vulnerability Analysis and Risk Assessment. ACSAC 2007.

[20] F. Kensing and K.H. Madsen. Generating Visions: Future Workshops and Metaphorical Design. In J. Greenbaum and M. Kyng, editors, Design at Work: Cooperative Design of Computer Systems. Lawrence Erlbaum, 1991.

[21] J. Landis and G. Koch. The Measurement of Observer Agreement for Categorical Data. Biometrics, 33, 1977.

[22] C. A. Le Dantec and W. K. Edwards. Designs on Dignity: Perceptions of Technology Among the Homeless. CHI 2008.

[23] C. Li, A. Raghunathan, and N. K. Jha. Hijacking an Insulin Pump: Security Attacks and Defenses for a Diabetes Therapy System. Healthcom 2011.

[24] J. K. Miller, B. Friedman, G. Jancke, and B. Gill. Value Tensions in Design: The Value Sensitive Design, Development, and Appropriation of a Corporation's Groupware System. GROUP 2007.

[25] L. I. Millett, B. Friedman, and E. Felten. Cookies and Web Browser Design: Toward Realizing Informed Consent Online. CHI 2001.

[26] L. P. Nathan, B. Friedman, P. V. Klasjna, S. K. Kane, and J. K. Miller. Envisioning Systemic Effects on Persons and Society Throughout Interactive System Design. DIS 2008.

[27] X. Ou. Ethnographic Fieldwork at a University IT Security Office. ACSAC 2013.

[28] N. Paul and T. Kohno. Security Risks, Low-tech User Interfaces, and Implantable Medical Devices: A Case Study with Insulin Pump Infusion Systems. HealthSec 2012.

[29] A. Raij, A. Ghosh, S. Kumar, M. Srivastava. Privacy Risks Emerging from the Adoption of Innocuous Wearable Sensors in the Mobile Environment. CHI 2011.

[30] K. B. Rasmussen, C. Castelluccia, T. S. Heydt-Benjamin, and S. Capkun. Proximity-Based Access Control for Implantable Medical Devices. ACM CCS 2009.

[31] M. Rushanan, C. Swanson, D. F. Kune, and A. D. Rubin. SoK: Security and Privacy in Implantable Medical Devices and Body Area Networks. IEEE S&P Symposium 2014.

[32] S. Schechter. Security That Is Meant To Be Skin Deep: Using Ultraviolet Micropigmentation To Store Emergency-Access Keys for Implantable Medical Devices. HealthSec 2010.

[33] S. Schechter, R. Dhamija, A. Ozment, and I. Fischer. The Emperor's New Security Indicators. IEEE S&P Symposium 2007.

[34] J. Sorber, M Shin, R. Peterson, C. Cornelius, S. Mare, A. Prasad, Z. Marois, E. Smithayer, and D. Kotz. An Amulet for Trustworthy Wearable mHealth. HotMobile 2012.

[35] Symposium on Usable Privacy and Security. http://cups.cs.cmu.edu/soups/.

[36] E. Troshynski, C. Lee, and P. Dourish. Accountabilities of Presence: Reframing Location-Based Systems. CHI 2008.

[37] K. K. Venkatasubramanian and S. K. S. Gupta. PSKA: Usable and Secure Key Agreement Scheme for Body Area Networks. IEEE Transactions on Information Technology in Biomedicine, 14(1), 2010.

[38] R. Wash. Folk Models of Home Computer Security. SOUPS 2010.

[39] A. Whitten and J. D. Tygar. Why Johnny Can't Encrypt: A Usability Evaluation of PGP 5.0. USENIX 1999.

[40] F. Xu, Z. Qin, C. C. Tan, B. Wang, and Q. Li. IMDGuard: Securing Implantable Medical Devices with the External Wearable Guardian. IEEE INFOCOM 2011.

[41] D. Yoo, M. Lake, T. Nilsen, M. E. Utter, R. Alsdorf, T. Bizimana, L. P. Nathan, M. Ring, E. J. Utter, R. F. Utter, and B. Friedman. Envisioning Across Generations: A Multi-Lifespan Information System for International Justice In Rwanda. CHI 2013.